

# THE ARCHITECTURE OF THE PICTURE SYSTEM

James F. Callan  
Evans & Sutherland Computer Corporation  
3 Research Road, Salt Lake City, Utah 84112

## Abstract

THE PICTURE SYSTEM is a self-contained, stand-alone line drawing system for presenting dynamically moving pictures of two- or three-dimensional objects. This highly interactive system can display smoothly changing true perspective views of 3-D objects in real time. Objects can be rotating, translating, and changing in scale. Individual sub-objects can assume independent or compound motion. Lines or parts of lines outside the chosen viewing window are clipped. All these transformations are performed digitally to avoid inaccuracy and range limitations.

In order to perform these calculations fast enough to show smooth motion, support high-level interaction, and also sustain flicker-free pictures of 11000 line segments or more, a new graphics system architecture has been developed. This paper will present each of the hardware components of the system and discuss what each contributes to the production of pictures.

THE PICTURE SYSTEM represents a significant new graphics system architecture in several respects. This paper will describe the system including data flow, the functions of each major component, and some of the advantages that accrue from the approaches taken.

Figure 1. shows a block diagram of THE PICTURE SYSTEM.

## The Picture Controller

Evans & Sutherland selected the DEC PDP-11 as the standard controller for THE PICTURE SYSTEM because of its popularity and because its architecture is well suited to supporting high-performance graphics systems. THE PICTURE SYSTEM requires no modifications to the PDP-11 at all. Therefore, it can be equipped with any standard peripherals, utilize any DEC supported system software, and can be put to general-purpose use when not executing a graphics program.

The UNIBUS which serves as the universal link connecting the PDP-11 processor, memory, and peripherals also serves as the connecting link to the controlling registers and graphic peripherals of THE PICTURE SYSTEM.

Prior to issuing any drawing commands, the Picture Controller passes parameters to the Picture Processor to indicate how subsequent coordinate data is to be interpreted and what transformations are to be made to that data. The data interpretation parameters indicate connectivity (i.e., how points are to be connected), and point of origin (i.e., whether the coordinate data is absolute or relative to preceding coordinates). It further states whether subsequent coordinate data is two-dimensional, three-dimensional, or homogeneous to provide a much

larger effective dynamic range than is normally available with 16-bit words.

After all these modes have been set up, the Picture Controller passes the actual coordinate data from its memory to the Picture Processor via a DMA interface. The program contains DMA setups for each block of coordinate data that can be sent without mode changes.

This process continues for the data describing all objects to be displayed, at which point the program waits for a clock interrupt indicating that the pre-specified amount of time allocated for a frame has elapsed. The program then jumps back to draw the next frame perhaps with a slightly different set of transformation parameters, to produce the illusion of smooth motion in the picture.

When there is a great deal of data to be drawn a mechanism is provided for organizing this data efficiently so that the entire frame can be passed to the Picture Processor within the time desired. This mechanism entails arranging into tables coordinate data which have identical mode settings so that the issuing of data will not have to be interrupted to reset any modes. To make this feasible on a large scale, the modes are made very general.

Thus it is possible to issue a command to draw an entire table of:

- a. connected lines from the current position;
- b. connected lines from a new position;
- c. unconnected lines from the current position;
- d. unconnected lines from a new position;
- e. dots; or
- f. characters.

Likewise it is possible to specify that the coordinates of points in a table are:

- a. all absolute;
- b. all relative to their predecessor; or
- c. all relative except the first, which is absolute.

Since coordinate data is being fetched from PDP-11 memory by the DMA, it is possible for the CPU to execute code simultaneously. This overlapping of operations substantially reduces the time required to process a frame.

The display program has access to a 120 cycle clock (which is actually in the Picture Processor) for setting refresh and update rates. This 120 cycle clock offers more choices than the usual 60 cycle clock, including the optimal refresh rate of 40 hz.

### The Picture Processor

The Picture Processor receives data sent by the Picture Controller in the form of two- or three-dimensional line endpoint coordinates. It operates on this data a point at a time, based on parameters passed previously by the Picture Controller, and deposits the processed data in the Refresh Buffer.

The Picture Processor starts by converting relative coordinate data to absolute. Conversion at this point, prior to any coordinate transformations, avoids accumulating the roundoff errors that may be introduced by transformation.

The Picture Processor then transforms the data by any pre-specified combination of translations, rotations, and changes in scale. Translation may be in any direction in three space; rotation may occur about any axis in three space; and scaling may be in any dimension in three space. These transformations are imparted by expressing the coordinates of each point as a four-component vector (by appending one or two pre-specified coordinates, if necessary, to those emanating from memory), expressing the desired transformation as a four-by-four matrix, and multiplying the two together to form a four-component vector representing the transformed point. The homogeneous nature of the representation introduces much flexibility and permits much larger transformation and scale values than would otherwise be possible.

Transformed data is next "clipped"; i.e., it is compared with a pre-specified viewing window, and lines or parts of lines outside this window are eliminated. Clipping is the most general way to perform windowing, and allows zooming into a displayed structure to any desired magnification.

In two-dimensions, the viewing window is a rectangular region in the drawing space. In three-dimensions it can be either a cubic region (for orthographic views) or a section of a pyramid whose apex is at the eye point (for perspective views). Clipping is performed with respect to all six surfaces of the viewing window.

At this point three-dimensional coordinates are converted to two dimensions by computing perspective or, if desired, orthographic views. An intensity value proportional to depth is also computed here for use in depth-cueing if desired.

The final stage in the Picture Processor's digital processing is a linear mapping of points from the objects' coordinate system into that of the Picture Display. The mapping is from the viewing window onto a program-specified region of the Picture Display called a viewport. The viewport may also be

thought of as six-sided, since in addition to the four geometric boundaries there are intensity maximum and minimum values which may be specified as well. When depth-cueing is used, lines at the front clipping plane of the 3D viewing window are assigned the maximum intensity and lines at back are assigned the minimum. Constant intensity is specified very simply by making the maximum and minimum the same.

The advantages of program-specified viewports are that they free the programmer from having to be concerned with the coordinate system of the output device in preparing his data, and that they provide a convenient way to show several views or pictures on the same output device.

The operations discussed above are performed serially on each point in the data base. That is, the coordinates of a point enter the Picture Processor, are converted to absolute homogeneous coordinates, multiplied by the transformation matrix, clipped to the window, put in depth-cued perspective, mapped into viewport coordinates, and deposited in the Refresh Buffer. A new point then enters the Picture Processor and is treated by the same operations. The total time required for all these processes ranges from 10 usec to 36 usec per point, depending on the complexity of the clipping operation, and averages about 20 usec. Thus the Picture Processor can process about 3300 points in 1/15 second, the largest frame time which supports smooth dynamics.

An important related feature of the Picture Processor is its ability to concatenate transformations. The matrices representing two or more simple transformations can be multiplied together to form a matrix representing a compound transformation, which is left in the Picture Processor for transforming points.

For temporary storage of matrices, a four-deep push-down stack is provided in the Picture Processor. Thus if a certain transformation must be compounded with another but will be needed subsequently, it may be pushed onto the stack before compounding and later popped off the stack for further usage. There is provision for overflowing the stack into PDP-11 memory in case a stack depth of four is not adequate for some program. The overflow has the same properties as the main transformation stack except that transfers to and from it are slower.

Data computed by the Picture Processor is normally deposited only in the Refresh Buffer, but can be read back into Picture Controller memory as well. This capability is useful for obtaining hard copy output and also for performing additional processing on pictorial output in software, such as slow time shading or hidden line removal. Results are available after the transformation stage, after the clipping stage, and after the viewport mapping stage. The viewport boundaries may exceed those of the Picture Display when output is to be returned to memory so that full 16-bit accuracy is maintained for cre-

ation of high-quality hard copy.

A final feature of the Picture Processor is its "hit test" capability, which allows the system to detect whether any part of a given picture element is within a program-specified region in the drawing space. Hit testing is used for implementing the pointing function with a data tablet, eliminating the need for a light pen.

#### The Refresh Buffer

Coordinate data output by the Picture Processor, still in digital form, is ready to be drawn on the Picture Display. It is not drawn immediately, however, but is instead deposited in a memory called the Refresh Buffer. Once deposited in the buffer, the data can be used to refresh the Picture Display any number of times before a new frame is ready for display. Therein lies the value of the Refresh Buffer and the reason why it is in the system: picture refresh rate and picture update rate need not be identical, so that each may be independently optimized. The refresh rate may be specified as any of 16 submultiples of 120 hz - usually 60, 40, or 30 hz. The generation of new frames need not proceed so fast, and so the update rate is usually lower than the refresh rate, typically 15 or 20 hz.

Because of the Refresh Buffer, the Picture Generator never has to sit idle while its input is being prepared. Therefore it can display many more lines in a given time than would be possible otherwise.

The buffer contains 8K 36-bit words, and can be optionally expanded to 16K 36-bit words. Each word can contain any of the following data types:

- a. a point to be moved to with the beam off
- b. a point to be moved to with the beam on
- c. a point where a dot is to appear
- d. three character codes
- e. mode settings for the Picture Generator or Character Generator.

In the first three types, the point is defined by a 12-bit "X" field, a 12-bit "Y" field, and an 8-bit intensity field. In the character words, there are three 8-bit character codes. In the last type, called "status", there are bits for indicating scope selection, dashed mode, blink mode, character height and width, and character orientation. These bits set up modes which remain in effect until another status word is encountered.

Data may be single-buffered, in which case the Refresh Buffer contains one frame of data which is overwritten by the next frame, or double-buffered, which entails splitting the Refresh Buffer into two areas so that a frame in one area remains intact for refresh purposes until a new frame is completed in

the other. The roles of the two areas are then reversed and a third frame is begun. Double-buffering is recommended when frames are small enough to fit in half of the Refresh Buffer, because single-buffered pictures can show any of a number of unpleasant effects that result from trying to display data that is simultaneously being updated.

A four-word area of the Refresh Buffer (two four-word areas in double-buffer mode) is devoted to a cursor so that a highly dynamic cursor can be maintained even if the update rate of the rest of the picture is slow.

#### Picture Generator and Picture Display

These two PICTURE SYSTEM components are discussed together since their specifications are so intertwined. The Picture Generator converts line endpoint and intensity coordinates, expressed digitally, to analog voltages for the Picture Display. The Picture Display contains a cathode-ray tube (CRT) whose electron beam sweeps out straight lines at any angle based on the analog voltages supplied by the Picture Generator. The beam briefly illuminates the phosphor coating of the CRT, which if repeated frequently enough results in a steady visible picture.

Lines may be solid or dashed, and may be made steady or blinking. Constant intensity of picture elements may be chosen from 256 levels. Lines are drawn at a constant rate which assures uniform brightness for the chosen intensity level. Depth cueing allows the intensity of lines to vary continuously with depth. In order to present a uniform variation in brightness, the intensity control of the Picture Display treats the Z coordinate data as the logarithm of the intensity to be displayed. The contrast control of the Picture Display is completely independent of the intensity control.

The following display rate examples assume a 30 hz refresh rate:

- a. 11,000 half inch lines connected end-to-end
- b. 1600 ten inch lines connected end-to-end
- c. 6600 uniformly distributed dots

The Picture Display has a 21" diagonal rectangular CRT with a quality viewing area of 10" x 10". Line endpoints may be centered on any point in a raster of 4096 x 4096 addressable points on the display surface. Line width and spot size do not exceed .020". Line endpoints match to within .020".

The standard phosphor of the Picture Display is P4, a white phosphor which fades very quickly after illumination so that no trail or smear is left by moving picture elements.

#### The Character Generator

When character words are read out of the Refresh Buffer, the Character Generator unpacks them into codes which access a read-

only memory containing character stroking data. The strokes are read out of the read-only memory one by one, multiplied by a pre-specified sizing parameter, and drawn by the Picture Generator on the Picture Display.

The read-only memory contains the 96-character extended ASCII character set, including upper and lower case alphabets, ten numerics, and thirty-four special characters.

There are eight character sizes available under program control ranging from .07" high in increments of .07" to .56" high on a 10" x 10" viewing area. The character width is also under program control with eight different widths selectable for each size. Characters may be displayed horizontally or in a 90° counter-clockwise orientation. Inter-character spacing is handled automatically by the Character Generator.

Such flexibilities as clipping, rotation, and continuous sizing cannot be imparted to characters generated by the Character Generator. When such operations are desired, the Picture Controller should produce characters just like other picture elements.

To place the Character Generator in front of the Picture Processor would severely and unnecessarily impact the picture update rate, since each individual character stroke would have to be transformed, clipped, etc.

#### The Tablet

The tablet serves as the standard, general-purpose graphic input device in THE PICTURE SYSTEM. Associated with the tablet is a pen whose coordinates are read by the Picture Controller. Normally a cursor is drawn on the Picture Display to indicate the position of the pen on the tablet.

The tablet is inherently well-suited to entry of precise positional information. With the aid of the Picture Processor's "hit test" capability discussed earlier, it can also be used conveniently for pointing to picture elements, eliminating the need for a light pen. An extremely versatile device, the tablet can also be used for command selection, entry of rate information, and even entry of alpha-numerics.

#### Options

The components described above comprise the so-called Standard Configuration of THE PICTURE SYSTEM. A number of options of a peripheral nature are also available such as keyboards, function switches, hard copy units and so forth. These units are interfaced directly to the UNIBUS. Additional Picture Displays and an expansion of the Refresh Buffer are also optional.

It is significant that there are no options within the central processing path of the system. All the processing features which E&S believes a high-performance system should have are standard in the system, not optional. The reason is that to make such a feature unplugable is to make it expensive.

It would have to have interfaces at either end to fit into the system; it would have to be separately packaged; and a great deal of care would have to be taken to make sure the system worked gracefully without it, hardware-wise, software-wise, diagnostic-wise and documentation-wise. Standardization is perhaps the most important reason why the system can be marketed for a reasonable price.

#### Conclusion

THE PICTURE SYSTEM's architecture represents a distillation of technology incorporated into earlier digital graphics systems as well as some new approaches. Widespread user acceptance indicates that the right set of features are included, and the relatively low system cost indicates that they are packaged in an efficient manner.

#### References

1. "Line Drawing System Model 1 System Reference Manual", Evans & Sutherland Computer Corporation, Salt Lake City, Utah January 1970.
2. "Line Drawing System Model 2 System Reference Manual", Evans & Sutherland Computer Corporation, Salt Lake City, Utah August 1971.
3. Callan, J. F., "Key Decisions in Designing THE PICTURE SYSTEM", SID Journal, Vol. II, No. 1, Jan/Feb 1974.
4. Newman, W. M., and Sproull, R. F., "Principles of Interactive Computer Graphics", McGraw Hill, 1973.

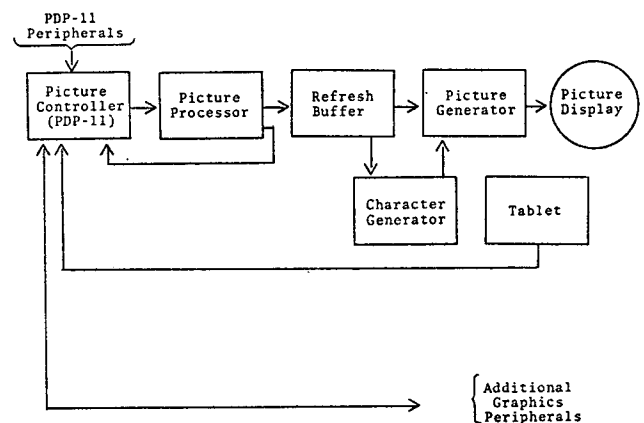


Figure 1.